Intel Developer Update is Intel's monthly online news magazine for developers. As the official publication of developer.intel.com, it brings hardware, software, and Web developers the latest information on Intel initiatives, technologies, platforms, and products.

Cover Story

Each month, we run a cover story on the most significant industry announcement, trend, or development for the month.

Featured Articles

Delivering in-depth reports on key platforms, products and technologies, our featured articles provide a monthly source of information on issues affecting developers. Be sure to check in every month for the latest developments driving the evolution of the industry.

Contact the Editor

To make *Intel Developer Update* a better information resource, we invite you to share your thoughts on what we've published or what you'd like to see covered. Comments are always welcome.

Archives

Our archives contain two groups of previously published articles. One group contains all the articles that appeared in *Platform Solutions News*, the earlier version of *Intel Developer Update*. The articles date from September 1997 through August 1999. The other group is set up to contain *Intel Developer Update* articles dating from the inaugural September/October 1999 issue.

Bookmarking

We advise against bookmarking article pages. They're accessible online only during the month the issue is live. Thereafter, they're removed to our archives. Instead, we suggest that you bookmark the PDF (Adobe* Portable Document Format) file versions of the articles. You'll find buttons for the PDF files labeled "print article" in the right navigation section of each article. A PDF for the entire issue is labeled "print magazine" and is located near top right side of the IDU home page.

Table of Contents

(Click on page number to jump to articles)

Note: Intel does not control the content on other company's Web sites or endorse other companies supplying products or services. Any links that take you off of Intel's Web site are provided for your convenience.

Cover Story

Delivering the Next-Generation Mobile PC Platform

Pankaj Kedia
Senior Marketing Manager
Mobile Platforms Group
Intel Corporation

Overview

Intel's new Mobile PC platform, based on the Mobile Intel® Pentium® III Processor-M and the Intel® 830 chipset, delivers significant advances in three areas, which can be thought of as the "three Ms" of next-generation mobile computing:

- Maximum Performance and Capability
- Mobility and Wireless Connectivity
- Multitude of Form Factors and Price Points

The Mobile Intel Pentium III Processor-M is the first mobile processor based on Intel's 0.13µ process technology, and it is Intel's highest performance mobile CPU. Together with the Intel 830 chipset, it delivers a host of new features, which enable the highest performance at the lowest power ever implemented in a mobile PC architecture.

The high-performance and low-power characteristics of the processor and chipset are available in a range of form factors and enable the mobile user to experience true mobility anytime, anywhere. These characteristics, along with the Intel® Hub Architecture in the Intel 830 chipset, make it easier for OEMs to enable wireless capabilities to the mobile user. Additionally, the scalable and stable nature of the platform makes it easier for OEMs to design and for IT to qualify and deploy.

Maximum Performance and Capability

The next-generation Mobile PC platform delivers significantly higher performance and capability driven by the new 0.13-micron process, the new Intel Mobile Pentium III Processor-M, and the Intel 830 chipset.

The Mobile Intel Pentium III Processor-M is the first in a series of products based on Intel's 0.13-micron process technology. Intel is the first company in the industry to bring 0.13-micron-based products to the market. This process technology is based on the industry's fastest transistors, which is enabled by short gate lengths and thin gate oxides. Additionally, high-performance copper interconnects and low-power SRAM memory make it possible to deliver up to 25 percent higher performance or up to 45 percent lower power. The Mobile Pentium III Processor-M features 44 million transistors, compared to a maximum transistor count of 3 million just five years ago. With such parameters, Intel expects the upcoming processors built on this 0.13-micron process to deliver significant performance headroom in the months ahead.

The new 0.13-micron process enables Intel to deliver processors up to 1.13 GHz at launch. The Mobile Intel Pentium III Processor-M includes other features that increase the performance of mobile PCs. These enhancements include doubling the capacity of the on-die L2 cache from 256 Kbytes to 512 Kbytes, enhanced data pre-fetch logic to pre-load data into the cache, and the introduction of 133-MHz FSB, a 33 percent speed increase for faster data movement.

The Mobile Intel Pentium III Processor-M at 1.13 GHz delivers about a 100 percent performance improvement compared to the 500 MHz Intel Pentium III processor from 18 months ago, and a 30 to 45 percent performance improvement compared to the Intel Pentium III processor at 1.0 GHz.

To deliver this performance without compromising battery life, Intel is also introducing the Enhanced Intel® SpeedStep™ technology. The original Intel SpeedStep™ technology ran the processor at full speed with the system running on AC power and shifted to a power-conserving lower speed on battery power. The Enhanced SpeedStep technology provides users the performance they need, whether the system is running on AC or battery power, based on usage model and application load requirements. The operating system dynamically adjusts the processor to its highest

clock rate for discrete intervals when required by applications such as video and graphics programs. Shifting speed dynamically for extremely short time intervals lets users experience maximum performance at all times, while optimizing battery life.

The Intel 830 chipset introduces a number of other performance enhancing features into the mobile PC segment:

- Intel® Hub Architecture provides a point-to-point interconnect for better I/O performance and improved concurrency with peripherals including IDE, audio, modem, and LAN devices.
- 133-MHz SDRAM provides 33 percent faster CPU access.
- 1 Gbyte of SDRAM memory doubles the amount of memory support.
- The new performance-optimized memory controller reduces memory latency by about 25 percent.

Greater Mobility by Design

In addition to delivering high performance and capability, the new processor and chipset from Intel enable enhanced mobility and wireless connectivity across a range of form factors.

Intel's 0.13µ process technology helps deliver significantly higher processor performance at power levels similar to today's mobile processors. This allows system manufacturers to deliver increasingly higher performance in thinner and lighter Mobile PCs.

The package technology used for the Mobile Intel Pentium III Processor-M, the micro Flip-Chip BGA and PGA, consumes less space and enables the design of thinner and lighter systems.

In addition, the Mobile Intel Pentium III Processor-M and the 830 chipset feature a new low-power mode called "Deeper Sleep." This new dynamic power management mode enables the system to deliver longer battery life by reducing the processor voltage below the minimum operating voltage while preserving the processor state. Deeper Sleep lowers the power as much as 60 percent more than today's "Deep Sleep" mode.

A typical mobile user running office applications is expected to be in the Active Power mode about 20 percent of the time, in Quick Start mode about 20 percent of the time, and in Deeper Sleep mode the remaining 60 percent of the time. Reducing power in any of these modes, and especially the Deeper Sleep mode, results in longer battery life.

Intel is also taking the lead in bringing power management features to other parts of the platform such as the chipset. In addition, the Intel 830 chipset features the integrated LAN connect interface, which saves cost and provides greater mobility.

Wireless Connectivity

Wireless connectivity is one of the defining attributes for greater mobility, and Intel's mobile PC building blocks are ideally suited for wireless clients in three important ways:

- The higher platform-level performance enabled by the processor and chipset support running multiple concurrent applications, in wired or wireless usage scenarios. This enables the mobile users to take their mobile experience on the road.
- Lower power and thermals of the processor and chipset, in combination with the smaller packages, enable state-of-the-art 1-inch "thin-n-light" notebooks with longer battery life. This enables mobile users to be away from their desks and experience true mobility.
- Intel Hub Architecture supports faster data rates and point-to-point serial I/O connectivity that is ideally suited for a variety of wireless connection options.

Multitude of Form Factors and Price Points

Intel has plans in place to offer the new Mobile Intel Pentium III processor to cover a range of form factors—Full Size Mobile PC, Thin-n-Light Mobile PC, Mini-Notebook, Sub-Notebook, and Tablet PC—over a period of time. In each one of these segments, Intel plans to deliver the highest performance at the lowest power based on the form factor in each segment.

The Intel 830 chipset represents a family of chipsets based on a single board architecture that can be deployed across multiple mobile platforms:

- The Intel 830 MP chipset, which is available now, includes an external graphics interconnect option that enables high-performance AGP 4X graphics.

In the second half of 2001, Intel plans to introduce two additional flavors of the 830 chipset:

- The Intel 830 M chipset with integrated graphics is designed for the mainstream performance mobile PC segment.
- The Intel 830 MG chipset provides integrated graphics performance for the value mobile platform segment.

The ability of a single chipset architecture to scale across the high-performance, mainstream, and value segments means that OEMs can consolidate testing and validation of multiple form factors on a single platform. The same board architecture can be used to support multiple form factors across a range of price points. The cross-platform scalability of the chipset also has advantages for enterprise IT managers, who can qualify multiple mobile PC platforms by performing the majority of the work on a single board.

Intel's fast ramp of the Mobile Intel Pentium III Processor-M roadmap is expected to make the processor available in multiple price points in a short period of time. Because it is designed to support mobile PC designs for 12 to 18 months, the 830 chipset architecture provides a stable environment that can help lower the total cost of ownership in corporate environments.

Summary

Mobile PC platforms based on the Mobile Intel Pentium III Processor-M and the Intel 830 chipset represent the next generation in Mobile PC computing. Based on the 0.13-micron process technology, the processor and the chipset deliver significant improvements in the benefits of ownership and reductions in the cost of ownership, making this the best time to buy a Mobile PC.

More Info

Technical articles, product briefs, white papers, and collateral information are available on Intel Web sites:

- Information for developers is available on the Intel Developer Site.
- Information for corporate IT and business users is available in the Business section of the Intel site.
- Information for home users is available in the Home Computing section of the Intel site.

Plan to attend the Mobile Computing Track at the Intel Developer Forum Conference Fall 2001 for more detailed information about the Intel's next-generation processors and chipsets.

The data cited in this article is based on CPU Test, 3D Winbench 2000 and MPEG Encoding, Video 2000 benchmark results with the following system configuration:

- Mobile Intel® Pentium® III Processor-M, 830M-based motherboard, 1.13-GHz Mobile processor with Enhanced Intel® SpeedStep™ Technology, 256 MB of PC133 SDRAM, 16-MB RDRAM LM, 20G HardDrive, 16x CD-ROM, Win2000, 10x7x16-bitMobile Intel® Pentium® III Processor 1-GHz, manufacturer system based on 440BX, 1-GHz Mobile processor-256k with Intel® SpeedStep™ Technology, 256 MB of PC100 SDRAM, 20G HardDrive, 16x CD-ROM, Win2000, ATI Rage Mobility*
- Mobile Intel® Pentium® III Processor 500-MHz, Manufacturer system based on 440BX, 500-MHz Mobile processor-256k with Intel® SpeedStep™ Technology, 128 MB of PC100 SDRAM, 20G HardDrive, 16x CD-ROM, Win98SE, ATI Rage Mobility*

Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel® products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance of Intel products, reference the Benchmark Limitations page or call (U.S.) 1-800-628-8686 or 1-916-356-3104.

Author Bio

Pankaj Kedia is responsible for Mobile Programs and Promotions Marketing in the Mobile Platforms Group at Intel. Pankaj has held various marketing, planning, and consulting roles in the high-technology industry for over 10 years, 4 of them with Intel. He holds an M.B.A. from Wharton, an M.S. from the University of Michigan, and a B.S. from the University of Roorkee in India.

Departments

Desktop

Intel® Desktop Boards: Your Choice for Windows* XP

Scott Bair
Senior Product Marketing Engineer
Desktop Platform Solutions Division - Marketing
Intel Corporation

Overview

The introduction of Microsoft Windows* XP consumer operating system and the availability of Intel® Pentium® 4 processors at mainstream price points is generating a powerful platform shift for the PC industry.

Because the foundation of a solid platform is the system's motherboard, PC builders and integrators have relied for years on Intel® Desktop Boards for performance, reliability, and consistent quality.

Selecting the right motherboard is an essential step in building the best possible foundation for the upcoming Windows XP transition. And choosing Intel Desktop Boards will help to minimize time-to-market, eliminate integration issues, and ensure the highest possible degree of customer satisfaction.

Intel provides validated desktop boards specifically optimized for Windows XP and backed by an industry-leading board optimization package that includes:

- Key product training and support during the Windows XP transition.
- Comprehensive testing including environmental parameters, drivers and BIOS, and peripheral compatibility.
- Windows Hardware Quality Labs (WHQL) certification.
- Value-added software from Intel and third-party vendors, distributed on CD.

Microsoft has announced a launch date of October 25, 2001 for Windows XP. To help system builders benefit from this momentum, Intel is introducing four new mainstream Pentium 4 processor-based desktop boards offering support for multiple memory types and Intel's Pentium 4 processor roadmap.

Together with the new Intel Desktop Board products, Intel's comprehensive board optimization program can help Intel customers make the most of what will likely be one of the most highly publicized product introductions in the history of the PC industry.

Time-to-Market Support

Systems based on Windows XP operating system and the Intel Pentium 4 processor feature the most advanced technologies ever implemented in a mainstream PC platform. Making the most of this technology leadership starts with technical support straight from the source.

The ability to obtain accurate answers to technical questions can be critically important to meeting product release deadlines and is especially important to benefit from the extensive marketing momentum anticipated with the Windows XP launch and the latest high-performance Pentium 4 processor from Intel.

Intel provides technical support for system builders in three critical areas:

- Updated driver support is available on the Web. By making beta drivers available, Intel is enabling integrators to perform evaluations prior to the Windows XP launch. Intel's support Web site provides answers to frequently asked questions, a list of known driver issues and an online form where driver issues can be reported.
- During periods of transition in the PC platform, timely answers to technical questions can be key to meeting product delivery deadlines. Intel is making answers to specific technical questions available to Intel product dealers through the Intel Customer Support Line.
- Intel provides training for Intel product dealers at biannual training events and focused technical training on an as-needed basis.

See the "More Info" section at the end of this article for the latest contact information.

Rigorous Testing

Testing is the proven way to ensure a solid platform. Intel's four-level test and validation program can help avoid integration issues, speed system development, and ensure the best possible customer experience. Comprehensive testing includes:

- Environmental testing—Intel conducts thermal, mechanical, and EMI hardware tests in the initial stages of board development. Tests include temperature and voltage margining, dynamic shock and vibration tests, signal quality, and clock jitter tests. This hardware testing establishes the foundation for a solid reference design that can be used across the motherboard industry.
- Component level testing—Intel tests individual software drivers for quality and robustness in conjunction with commonly used applications in multiple operating environments. In addition, the Intel developed BIOS, undergoes rigorous testing, including focused suspend-and-resume (S3) power-management testing. Comprehensive component testing provides the basis for platform stability.
- Platform compatibility testing—Following component level tests, Intel conducts an extensive number of platform interoperability tests with over 350 commercially available add-in peripherals on 10 different operating systems with over 70 commonly used software applications. Tested components include network adapters, video cards, CD/DVD-ROM devices, USB peripherals, and PCI add-in cards. This testing results in a rock-solid platform that works with the latest technology, including Windows XP and the Intel Pentium 4 processor.
- Motherboard certifications—Following our internal tests, Intel submits its desktop boards to various industry and third- party certifications, including the Windows Hardware Quality Labs (WHQL). Designing platforms with certified components enables system builders to minimize integration issues. In key platform transitions, like the release of Windows XP, having a well-tested and certified desktop board is the key to providing time-to-market systems.

Intel® Desktop Board CD

Intel Desktop Board software CD provides a suite of utilities and applications that add value, speed system integration, and support product differentiation.
Intel® Express Installer is a convenient utility that simplifies installation of the software and drivers included with Intel desktop boards. This convenient utility takes the guesswork out of selecting the correct drivers by automatically recommending the most appropriate drivers based on each specific system configuration.

Security protection is vital to Internet-based e-Commerce, personal and financial data stored on PCs, and confidential e-mail shared over public networks. Intel now offers a full suite of third-party software to help make new PCs Internet-ready.

Today, even systems used in a small business are expected to be multimedia-ready. With the proliferation of DVD and CD-RW devices, users must be equipped with the essential software components they need. Intel desktop boards include a comprehensive selection of applications and utilities designed to enable the best possible user experience.

The Intel desktop board CD includes:

- Norton Internet Security*—includes Norton Antivirus*, Norton Personal Firewall*, Norton Privacy Control*, Norton AntiVirus*, and Norton Parental Control*.
- RealPlayer* Basic—optimum performance with Internet streaming media.
- RealJukebox* Basic—record and manage MP3 music files.
- SoundMAX* with SPX—3D positional sound with advanced audio rendering technology.
- NTI CD-Maker 2000*—create audio and data CDs.
- Intel® Rapid BIOS Boot—speeds Power On Self Test (POST) routine in speed system boot with Windows XP to cut boot time to as little as 20 seconds.
- Intel® Active Monitor—alerts users of thermal issues or fan failures.

Summary

The introduction of Windows XP will likely represent one of the most highly publicized transitions in the history of the PC industry.

Intel, the leader in desktop PC technology, is introducing four new desktop boards at mainstream price points designed to optimize the performance of Intel Pentium 4 processors and chipsets, while helping system builders capture the momentum of the Windows XP OS roll-out.

The selection of the right desktop board is the foundation for fast time-to-market development and satisfied customers. Intel desktop boards provide a full range of optimizations designed to increase performance, facilitate integration, and ensure trouble-free user experiences:

- Thorough technical support and training for system builders.
- Rigorous testing and validation of hardware, drivers, and BIOS, with compatibility and interoperability testing, and WHQL certification.
- A full suite of software, including utilities for easier integration and applications for Internet security and multimedia.

By choosing Intel desktop boards, combined with Intel's value-added support, system builders can get to market faster with high-quality products to get the most from the launch of Windows XP.

More Info

For the most current information visit the Intel desktop boards Web site.

For driver support, a list of known driver issues, and to report issues, visit the Drivers Evaluation Program Web site.

Visit the Intel desktop board software site for the latest information on Intel's suite of software utilities and applications. For technical support on third-party software supplied with Intel desktop boards, contact the appropriate third-party supplier directly.

For technical support on Intel desktop boards, or additional information, call (800) 628-8686.

Author Bio

Scott Bair has spent three years at Intel where he has been involved in programs and products including Intel® AnswerExpress℠ Support Suite, Intel Active Monitor, Intel Express Installer, Intel Desktop Board CD. He holds M.B.A. and B.S.E.E. degrees from Brigham Young University.

Initiatives and Technologies

Compliance Testing for USB 2.0 Devices

Doug Gemmill
Program Manager
Platform Industry Enabling
Intel Corporation

Overview

The Universal Serial Bus Implementers Forum (USB-IF) has taken the position that all USB products must function together without problems. Each class of USB product must pass specific testing, and the devices must also pass testing for interoperability. This also means that low-speed, full-speed, and high-speed devices must function together in the same USB tree. Full-speed devices must be able to work properly in a high-speed environment, and vice versa.

In the beginning, compliance testing for USB devices was a rather cursory process of plugging in simply to see if the device enumerated. Since then, compliance testing has evolved into a robust compliance program with detailed test matrices and multiple tests, including tests for interoperability. Products that pass today's compliance testing are allowed to display the new, high-speed logo for USB 2.0 a logo which gives users the confidence that all USB devices will work together, whether low-speed, full-speed, or high-speed.

Today's USB test procedures (developed by the USB Implementers Forum, along with Intel and other companies) check these capabilities:

- Electrical characteristics
- Device firmware (if applicable)
- General interoperability between the device being tested and other known-good devices, hubs, and hosts

Where to Test

Before a manufacturer can display the new, trademarked, high-speed USB-IF logo, the manufacturer’s device must pass compliance testing. There are two options for testing:

- Test at a USB-IF sponsored plugfest, which are offered several times a year and are free for USB-IF members.
- Test at an independent lab.

Plugfests offer an opportunity for vendors to see how well their products work together. However, independent test labs have several advantages over plugfests. First, plugfests occur only a few times a year. Plugfests also draw a vendor’s competition, so marketing confidentiality and other issues can arise. At independent labs, testing remains confidential. Independent labs may also offer detailed debugging services for USB devices.

Cable and connector testing is only done at independent test labs due to their extremely involved nature. The testing for cables, connectors, and cable assemblies take multiple days’ work and highly specialized testing equipment and test chambers.

Test labs that can certify USB 2.0 devices are listed at the USB developer's site. The USB-IF hopes to begin certifying independent test labs to perform testing on all USB devices by August 2001.

Pretest

During a plugfest, USB device testing begins with a pretest. The pretest is performed on a PC system or other host that does not have a USB device driver installed. The pretest follows this general process:

1. Connect the product to the host. Required events: the device enumerates cleanly; the device installs without error; no “blue screens” appear; and the user does not have to restart the host after the driver is installed.
2. Do an asynchronous disconnect of the device from the host.
3. Using a different port, reconnect the device to the host. The device must enumerate cleanly and be ready to use.
4. Do another asynchronous disconnect of the device from the host.
5. Using a third port (if available), reconnect the device to the host. Again, the device must enumerate cleanly and be ready to use.

If the device successfully passes the pretest, it is ready for compliance testing.

Electrical Test

The four electrical tests performed on full-speed and low-speed USB devices are described in the test procedures available on the compliance page of the USB-IF Web site. The four basic electrical tests consist of:

- *Signal quality test* makes sure the device meets specification for voltage levels and crossover times. This test is performed five hubs deep, which is the highest number of hubs allowed by USB 2.0. The hubs used in these tests have only marginal signal quality, so the device under test (DUT) is stressed as much as possible.
- *Inrush current measurement* measures the amount of current the DUT consumes when it is first connected to the system.
- *Power consumption measurement* makes sure the DUT is not consuming more power than allowed. Power consumption is measured and verified for unconfigured, configured, operating, and suspended states.
- *Backvoltage test* makes sure the DUT is not applying current on D+ and/or D- without first having VBUS present.

Interoperability

Interoperability testing makes sure the DUT works correctly in a fully loaded system of five hubs and other devices. At a plugfest, these other devices consist of one mass-storage device, one mouse, and one isochronous device. Interoperability tests check four main capabilities:

- *Accurate enumeration* at the end of the tree, five hubs deep, behind both an OHCI (open host controller interface) and a UHCI (universal host controller interface) host controller. Basically, these tests make sure the DUT functions properly. They also make sure the DUT does not cause other devices in the tree to fail.
- *S3 (suspend-to-RAM) state operation and return from S3.* For example, testing makes sure that the host returns to S1 when devices with remote-wake (wake-on-LAN, wake-on-ring) capability receive the appropriate remote signal.
- *Functionality after reconnect* to another location in the USB tree.
- *Shutdown, reboot, and cold and warm restart cycles.*

High-Speed USB

High-speed devices undergo additional tests at a plugfest. Once these DUTs pass all full-speed tests, they are subjected to both high-speed electrical and interoperability testing. Because of the nature of high-speed USB, these tests are much more involved than the testing for full-speed devices.

In high-speed tests, the host puts the DUT in a test mode. The DUT cannot exit this test mode until power is recycled to the DUT. This isolates the DUT from the host and allows the DUT to be specifically and rigorously checked. The 2.0 USB specification (available on the USB-IF site) describes the test modes that high-speed USB devices must support.

High-speed testing procedures are currently being documented. These procedures will be posted in the developer's area of the USB-IF site, along with the 2.0 USB specification and other documentation.

Cables, Connectors and Cable Assemblies

A successful 2.0 USB initiative depends on quality cable assemblies for connecting devices to host computers and hubs. So the test requirements for cables and connectors have been significantly enhanced. All cable assembly and/or connector testing is done in specially certified test labs. These test labs perform all required tests, including mixed-flowing gas, attenuation, shielding, shock, and durability. The mixed-flowing gas test, for example, involves placing a connector or cable assembly unmated into a special air-tight chamber while certain corrosive gases are introduced to the chamber and circulated around the connectors (or cable assemblies) for five days duration, then tested for low-level circuit resistance, then placed back into the chamber for five more days, only mated into its corresponding receptacle, followed again by low-level circuit resistance level testing again. Similar testing is done in a high-humidity chamber (connectors are tested both unmated and mated) to ensure the materials used will withstand long periods of use in sometimes rather intensive conditions.

During the tests, a pigtail is attached to the connector, and the receptacle is mounted on a test jig. The connector and receptacle are then subjected to several days of testing, as described in the connector test matrix. The test matrix is provided on the USB-IF site. The test fixture used in this procedure is also supplied free to members by the USB-IF to help reduce vendor testing costs and to increase the consistency of test results.

Also, devices that connect to the USB must be constructed with standard USB series-B or mini-B receptacles that have already undergone and passed compliance testing at a certified independent cable and connector test lab. These receptacles are listed on the USB-IF Integrators List on the USB-IF Web site.

Cable assemblies must be constructed with cables that have the Underwriter's Laboratories seal. Cables with this seal are assumed to be fully functional, and are not specifically tested. Cable assemblies must also be constructed with connectors that have already passed testing. The USB cable and connector specification can be downloaded from the USB site.

Devices that connect to a USB must be constructed with standard USB receptacles that have already passed testing. These receptacles are listed on the USB-IF Integrators List on the USB-IF Web site. Any 2.0 USB-compliant product that ships with a cable assembly must use a cable assembly from the Integrators List.

Similarity Tests

The testing of one product may allow another, very similar product to be added to the Integrators List without being individually tested. This is called certification by similarity. For example, a cable assembly might be eligible for certification by similarity. If the components of the cable assembly have already passed test on the old product, they are already certified, with some limitations, to perform properly in the new cable assembly.

For example, a cable assembly may be 4.5 meters long and use 28 AWG cable. In this case, the connectors on the cable assembly have already passed USB 2.0 testing. The manufacturer of this cable assembly can apply for certification by similarity for another cable assembly, as long as these conditions are met:

- The connectors used on the new cable assembly are the same as on the previously tested cable assembly.
- The same manufacturing process is used for the new cable assembly.
- The cables are a shorter length cable (for example, 3 meters), or the cables are thicker-gauge wire (such as 24 AWG).

Products with significant differences must still be individually tested before being certified. In addition, information about the new cable assembly must be sent in to the USB-IF office before the product can bear the high-speed USB logo. Once the new product is listed on the Integrator's list, the manufacturer can use the high-speed logo on the new assembly.

Summary

A proliferation of USB ports, hubs, and devices is entering today's market. Consumers need confidence that the USB devices they buy today will work with—and in—both new and existing systems. The new high-speed 2.0 USB logo gives users the confidence that all USB devices will work together, whether low-speed, full-speed, or high-speed.

Compliance testing for USB devices, cables, and connectors is available at independent test labs. Devices, hubs, and hosts may also be tested at USB-IF sponsored plugfests. Check the USB developer's site for the next plugfest schedule, lists of recommended labs, test specifications, and other necessary compliance information. Developers should insist that all USB devices are tested and certified for USB 2.0, as described on the USB-IF site.

More Info

Additional information about the USB 2.0 compliance program, including the new, trademarked high-speed logo, is available in the Intel Developer's Update article titled "USB 2.0 Is out of the Starting Gate."

The USB-IF developer's site includes lists of recommended test labs, the 2.0 USB specification, implementation information, and more. The test procedures for low- and full-speed USB devices can be downloaded in PDF format from the USB site. The site's extensive FAQ also includes information about compliance issues. And, testing techniques for low- and full-speed devices are described in a number of white papers that cover cyclic redundancy checks, HID parser error checking, and other topics. Procedures for high-speed compliance testing should be posted on the site in October of this year.

The site also includes device class specifications, the UHCI specification, version 1.1, and the OHCI specification, version 1.0a. Independent test labs that are already certified to test USB 2.0 products are listed on the USB site.

Author Bio

Doug Gemmill is the program manager for USB in Intel's Technology Initiatives Marketing group. He currently serves as compliance committee chair for the USB-IF.

Before joining Intel, Doug worked at Synopsis Logic Modeling, where he held various positions in technical marketing. He earned his bachelor's degree at Biola University.

Networking & Communications

Innovative Thermal Solutions for a High-Density Network Appliance

Mark Summers
Senior Thermal Mechanical Engineer
Embedded Intel Architecture Division
Intel Corporation

Javier Leija
Senior Thermal Mechanical Engineer
Embedded Intel Architecture Division
Intel Corporation

Overview

High-performance Internet infrastructure appliances can be used for security encryption and decryption, network attached storage (NAS), virtual private networks (VPN), and Voice over IP. The increasing requirements for high density in such devices underscore the value of packaging as much processing power as possible into a small footprint, 1U form factor. A flexible platform architecture that implements CPU and network I/O functions on separate boards has the additional advantage of enabling a single embedded Intel® Architecture design to support a variety of products with custom I/O configurations.

Design teams at the Embedded Intel Architecture Division (EID) create proof of concept reference designs that enable customers to reduce time-to-market. Working on such designs also makes it possible for Intel to help resolve some potentially difficult design challenges.

In the case of the performance communications reference design, which features dual 866-MHz Pentium® III processors, four hard disk drives, Gigabit Ethernet I/O, and other components in a 1U chassis, thermal management is one of the foremost design challenges. The proof of concept design involved the use of an innovative high-reliability thermal solution that was tested by advanced computational fluid dynamic and thermal modeling tools.

Platform Architecture

The Intel design team was presented with a list of requirements including the implementation of dual Pentium III processors at 866 MHz with a thermal design power (TDP) specification of 26 Watts per CPU, PC133 SDRAM with error correction code (ECC), 64-bit/66-MHz PCI expansion capability, four Ultra160 SCSI hard drives rated at 24 Watts, Gigabit Ethernet with an optical interface, three 10/100 Ethernet ports, and a PCI mezzanine card (PMC) for additional I/O expandability—all packed into a 1U chassis with an internal power supply. Due to the requirement for PC133 SDRAM and support for dual Pentium III processors, the design team selected the ServerWorks ServerSet* III HE-SL northbridge and OSB4 southbridge (see Figure 1). The 64-bit PCI bridge employs a proprietary ServerWorks 16-bit/266 InterModule Bus (IMB).

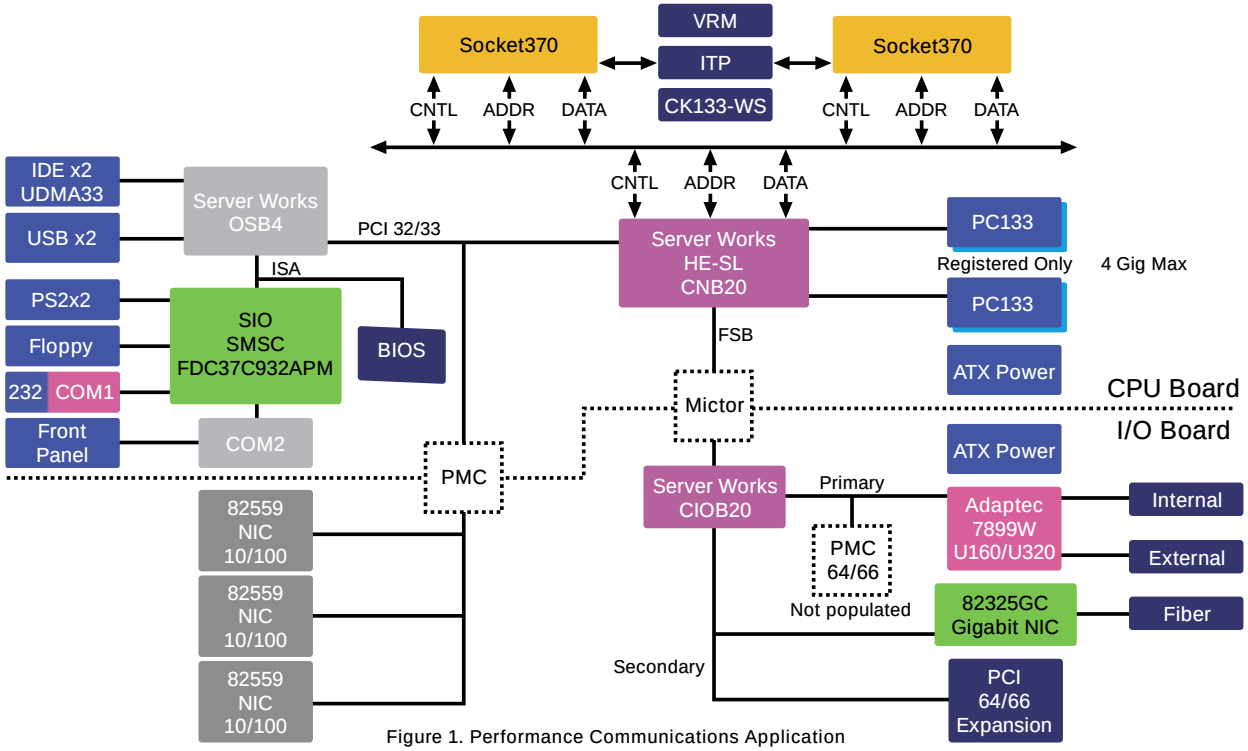

Figure 1. Performance Communications Application

CPU functions and I/O are implemented on separate system boards to accommodate the requirement for the smallest possible X-Y footprint and to provide design flexibility. The Intel design team worked with ServerWorks, Inc. to develop topologies and termination to enable the IMB to span across the boards.

Taking the 16-bit bus across the boards has the advantage of keeping the wide fast buses, including the front-side bus and the dual-channel memory bus, entirely on the local boards. The z-axis space constraint imposed by the two-board solution required the use of an angled DIMM connector for system memory, which in turn necessitated simulation to validate topologies and routing. The power supply and chassis were modified OEM units.

The PMC enables custom video support; however, the communications reference design is headless, with COM1 provided for command and control. The IMB and PMC enable designers to use a single CPU board to support multiple product configurations that can be customized with varying I/O configurations on the I/O board. The I/O board features dual independent PCI 66 MHz, 64-bit channels that support peer-to-peer writes. The Gigabit Ethernet NIC and SCSI are implemented on separate channels to enable maximum throughput. The onboard voltage regulation module (VRM) requires no external cooling.

Thermal Challenge

The thermal solution was one of the major challenges in the design of this high-density solution. The primary constraint was packaging dual Pentium III processors in FC-PGA form factor and with a maximum junction temperature (Tj) of 75°C in a small footprint, 1U OEM chassis, which permits a maximum z-height of 1.72 inches (43.6 mm).

The two-board (CPU board and I/O board) design employed a stacked board configuration. The presence of the second board meant that the design team had even less than the full 1U z-height available, further compounding the thermal challenge.

All components added for thermal management and cooling are selected to support optimal system reliability over a typical five-to-seven-year life and are designed to support maintenance and upgrades. The thermal solution was designed to enable operation in an ambient (outside) temperature environment of 50°C. Analysis showed that all of the electronic components within the chassis would raise the internal chassis temperature by 15°C. This left a thermal design window of 10°C. It is illustrative to compare this with a desktop PC, where thermal design window is at least 30°C.

Thermal Solution

The thermal design solution incorporates a two-compartment package (see Figure 2). The forward compartment contains four SCSI hard drives, and the rear compartment contains the CPU and I/O boards. A fan wall divides the two compartments, with five system fans to pull ambient air into the front compartment, into the rear compartment, and then out through exhaust vents and fans. The five system fans that come standard with the OEM chassis are rated at 8.8 CFM and have a seven-year design life.

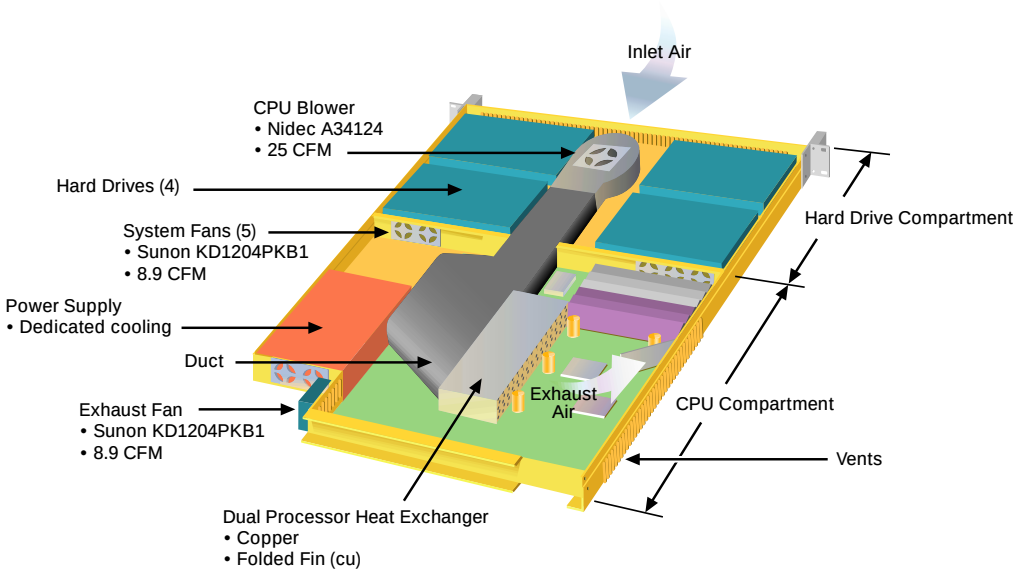

Figure 2. 3D CAD image of the Communications Appliance Reference Design thermal management solution.

Thermal modeling showed that air passing over the four hard drives in the front compartment would experience a 15°C temperature rise. Because this exceeds the allowable thermal window, an additional solution was required.

The solution featured a blower to direct additional outside air through a dual-channel duct, bypassing the front compartment, to a heat exchanger designed for the dual Pentium III processors. The blower is rated at 25 cubic feet per minute and has a design life expectancy of 11 years. It includes a tachometer for system monitoring. The duct is a plastic molded device designed to be easy to manufacture. Its internal channels are designed to equalize the airflow to each of the two processors. The heat exchanger for the CPUs is a single copper folded-fin exchanger with excellent heat dissipation properties.

The blower, duct, and heat exchanger combination increases the allowable thermal design window up to 25°C (75°C maximum Tj for the processors minus the 50°C ambient air temperature).

Computerized Simulation

The design team used computational fluid dynamic and thermal modeling tools to simulate the behavior of components including the 1U chassis, the network appliance platform electronics, and the overall operating environment. This simulation analysis had five goals:

- Calculate the junction temperatures of critical components.
- Calculate the maximum pressure inside the chassis to ensure optimum fan function.
- Compare the results to component specifications.
- Optimize components to achieve the best thermal solution.
- Compare the simulations to empirical results.

The thermal modeling results could be used to optimize the configuration of the chassis. The results showed that under worst-case conditions the maximum Tj values for the Pentium III processors (up to 30 Watt power dissipation) is 72°C, which is below the maximum allowable Tj of 75°C.

The simulation model was also used to simulate the failure of one of the system fans. Results showed that even in the event of a failure, the remaining four fans could provide adequate cooling within specified parameters.

Separate the 1U chassis into two compartments and model them separately.

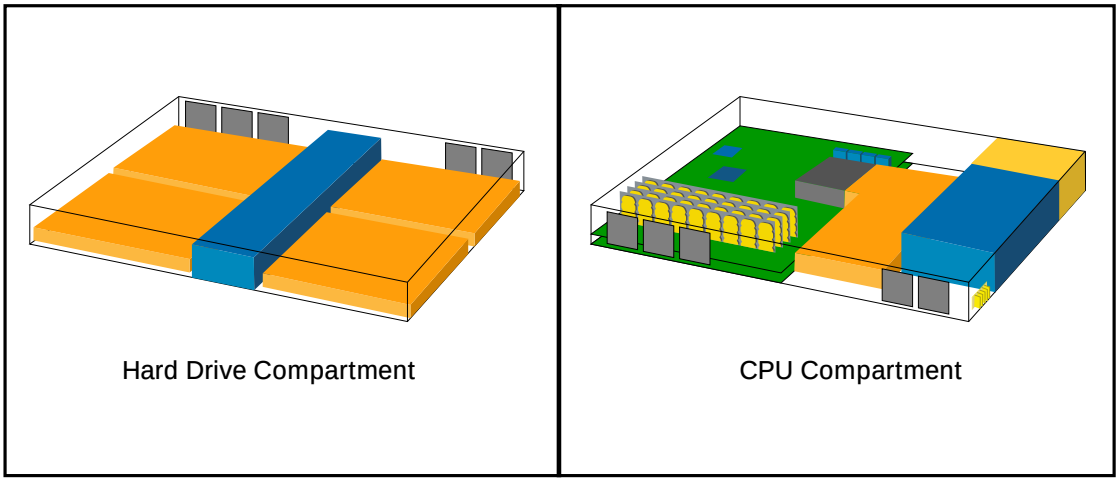

Figure 3. 3D thermal models used to optimize the thermal management solution for the Communications Application Reference Design.

Summary

This thermal solution meets design constraints by cooling the dual Pentium III processors and other critical components within the 1U chassis, while maintaining system reliability. The system is maintainable and upgradable without the use of specialized tools. The design is optimized for an ambient temperature of 50°C by raising the thermal design window from 10°C to 25°C. Thermal analysis tools were used to optimize the design and demonstrate that it can adequately cool dual Pentium III processors in a FC-PGA form factor rated up to 30 Watts power dissipation per processor.

This proof-of-concept helps demonstrate the value of a high-density design featuring dual Pentium III processors in a flexible dual board platform configuration that abstracts embedded Intel Architecture CPU functions from I/O. This enables manufacturers to support a variety of fast time-to-market product designs based on a single CPU board. Developers can use the CPU board only and add I/O capabilities as required with a custom I/O board. The proof of concept design can also be repackaged to create a flexible small office/home office Internet appliance.

More Info

The Intel® Performance Communications Reference Design described in this article is available for download at no cost on the Intel Developer's Web site. It includes schematics, design rules, and a user's manual.

An Interactive Thermal Design Demo is available on the Embedded Intel Architecture in Communications Web site.

Author Bio

Mark Summers is a senior mechanical engineer in Intel's Embedded Intel Architecture Division. Mark has 16 years of experience in system and component electronics packaging, with a specialization in rapid prototyping and severe environment systems packaging and analysis. He graduated from Arizona State University with a B.S.M.E., and has completed graduate level course work. Previously, Mark spent 15 years at Motorola. He currently holds 14 U.S. patents and has written several technical papers.

Javier Leija is a mechanical engineer with six years experience in system and component electronic packaging with a specialization in mechanical and thermal analysis. Javier graduated from Arizona State University with a B.S.M.E., and has completed graduate level course work. Javier currently works as a senior thermal mechanical engineer for Intel Corporation's Embedded Intel Architecture Division. Javier has one issued utility patent, several patents currently pending, and has presented several papers at national conferences.

Servers

Enabling a Client-Edge-Server Internet Model

Michael Condry
Director
Network Edge Technologies
Intel Corporation

Overview

Today's Internet operates in a client-server model for data transfer. However, the model is less than inadequate. The content being transferred is increasingly complex and exotic. In addition, clients are much more diverse. In short, it's more and more difficult to efficiently deliver complex, varied, and exotic content to increasingly diverse client devices.

There's a temptation to continue as we have done in the past, to simply create new servers and clients—new "boxes"—to handle each new type of data. For example, we could create streaming servers to get streaming data out to clients. However, it's a temptation we can resist because the structure of the Internet offers a more efficient and exciting alternative than building yet another box. Network edge architecture is the new architecture that defines the way the content is provided by distributing the required processing.

Today's Model

In the conceptual client-server model, the client requests information from the server, and the server sends information back in the format requested by the client. This content delivery system was once considered ideal since one need only concentrate on processing at the "ends." However, reality has demanded more, and in response new services have emerged. Caches near the client and on edge-delivery points make operations faster and lighten the network load. Gateways (such as those for wireless devices) help diverse groups of clients connect. And we now have Content Delivery Networks (CDNs) to better organize the vast array of Internet content for the client.

The problem is that these services have grown independent of one another, making extensions and new services vendor-specific and hard to manage. Developing network services has become difficult and expensive, prohibitively so for many independent software vendors (ISVs) and for service providers as well. With the evolution of diverse client devices, content providers are forced to offer data in every necessary format for every device, personal data assistants (PDAs), personal computers (PCs), cell phones, laptop computers, e-book readers, and others. Servers are forced to handle differences between the formats available and the formats the client can accept. The development and operation costs of incompatibly architected services have weakened market growth.

Network Edge Technology

A compatible architecture provides a means to address the complexities that have grown for "in the network" services. The term "network edge" is usually used to describe a physical or network boundary of an entity in the Internet. For example, network edges could be the set of access points of an Internet service provider (such as a DSLAM or a cable head-end) or the intranet of an enterprise network. Yet, "network edge" can also mean a distinguished point of content delivery, such as the boundary of a country, a point at which you want data translated into an appropriate language or format. Thus, the network edge is the point where content processing can occur in the path of content delivery.

In this context, network edge technology, is a collection of technologies for performing part of the content delivery processes for transferring data between clients and servers. Its architecture is a unified approach to developing software and hardware for performing content delivery functions.

Network edge technology architecture allows for operations such as data transcoding. Consequently, content providers can offer data in a single, standard format regardless of the client's device, and network edge technology makes the presentation suited to the device. For example, a content provider could store content in a single format, regardless of the client's device, and the network edge technology would reformat it for the receiving device. Rule-based service running at the edge makes the appropriate data transformations based on standardized protocols for various devices.

Consider an example of a movie delivery. With network edge technology, consumers see the movie in the format appropriate to their devices and their personal tastes. The content provider no longer has to worry about the device on which the movie will be viewed, since the burden of data transformation is removed from content providers and servers. Instead, data transformation is consolidated at the network edge. For quality assurance purposes, the content provider controls the transcoding options so that the content is not transformed in an unacceptable manner.

Personalizing Requests

Today, customer information and client device interfaces must be managed on the client device and/or content server, requiring the customer to re-enter data and the server to cope with the client diversity issues, such as cell phone vs. personal computer characteristics.

The transfer and reformulation of content from provider to client happens in several steps. First, the client makes the request. The request is amended (personalized) by known information retained in the server with additional needed data being entered from the client. For example, the client may be using a cell phone to access the net. This requires that the server send back only content appropriate to that cell phone interface and progress through voice menus to obtain customer specifics. In the network edge technolgoy environment, the edge device amends the client's request to specify the appropriate cell phone interface. Then the edge device communicates with the content servers to fill in customer data simplifying customer interaction then finally amends the request to the servers to obtain data that is reformulated to fit the client's device. The appropriate content is then downloaded to the device in a manner suitable for the client device.

The network edge technology device sits between the client and the content servers and serves as an agent for both the customer and the content provider. When the client sends out a request for content, network edge technology amends the request based on a set of predefined rules according to the customer’s information. Network edge technology then sends the amended request on to the server, possibly with some exchange to the content server as indicated previously. The server does not have to amend the request or reformulate the content. Any transformation of content that is needed will happen at the network edge device using standard protocols. This simplifies the information required from the content provider to address different formats needed for different devices.

Network Edge Technology Components

Network edge technology is a consistent, rule-based service environment. It provides a collection of standardized protocols and application programming interfaces (APIs). These protocols and APIs let a control system (for example, the extension of a proxy) take a request, see what has to be done to that request, get the content needed, do the transformation, and send that content back to the client.

Network edge technology includes:

- A rule-based engine
- A rule language, which is currently being defined
- A set of service “callout” protocols
- Standards for content representation (metadata) and user/device profile
- Policy framework
- Security framework
- An open-source reference implementation of the basic building blocks

The rule system defines the functions and rules for operations performed by both the client and the content provider. With rule-based service, content requests are automatically personalized for consumers, based on their device or user preferences. Services are written as “plug-ins” to the base framework or “callout” services where content is redirected for the service to be performed. The returned content is also automatically resized and reshaped to their devices. This extends consumers’ ability to access content and have an enjoyable Internet experience.

Content providers can still put restrictions on how content is delivered. For example, restrictions could be used to make sure content is not degraded by poor or inappropriate data transformations.

IETF/OPES

The Internet Engineering Task Force (IETF) has created a new working group for network edge technology. This group is called OPES (pronounced *oh-pez*) for Open Pluggable Edge Services. Its purpose is to build a general architecture for content delivery at the network edge. OPES will also create the protocols needed to make that architecture work. This is the framework that will enable the deployment of network edge services.

Although the IETF/OPES working group is just getting started, industry is already showing strong interest. Preliminary and interim working group meetings have drawn active participation from over 20 companies, and requirement documents are already being created. The working group itself already includes over 100 active members, including some very active representatives from across the network, caching, and computer industries.

Some consumers and vendors worry about defining ”middle of the net services” standard or not. These concerns stem from operations using transparent proxies that, without permission, insert ads and perform other intrusive tasks. However, network edge technology services are explicitly requested and explicitly addressed in the IP arena. The services are only those requested by the client or content provider.

Intel's Input

OPES is an open-systems strategy led by Intel and others, but not Intel-owned. The strategy is focused on an open-source solution that has multiple-chip architectures to create the most efficient edge design. Vendors can design their software and hardware with any appropriate processors or software base.

Intel does offer developers an array of chip architectures through its IA-32, Itanium™ (IPF), and Internet Exchange Architecture (IXA) families of processors. Because each family has its strengths in specific roles, a developer could create an edge service that incorporates different types of processors.

Services such as SSL would run faster on IPF than on regular IA-32 processors. Data stream conversion tasks are also likely to run faster on IPF than on IA-32-based machines. In contrast, the IA-32 processors would effectively handle web assembly, wireless transformations, and other functions. IXA also offers compatible control components. As with other technologies, multiple Intel® architectures used in edge technology can give OEMs and service providers a key advantage.

Summary

Network edge technology is the result of a growing shift from today's client-server model to a more efficient, client-edge-server model. When content-level services are performed at the network edge, content providers can offer personalized, scalable services. For example, caching companies could extend the capabilities of their hardware by adding software. In fact, caching companies may be the first to add edge services to the traditional client-server model.

With network edge architecture, the edge plays a role that is currently forced on both client and server. And, because the network edge is the client's point of entry, the market for edge technology is much greater than the market for servers. In short, the network edge has become the critical area of growth for the Internet.

Developers should follow the IETF/OPES group’s work closely. Developers should also contribute actively to the on-going work to define the edge architecture and network edge technology interface and protocols. The Intel Developer Forum Conference Fall ’01 being held in August, will offer several sessions on network edge technology.

More Info

Information about edge services is available online at the Stardust site.

Developers can also check out the dedicated site for the OPES working group. This site includes a relatively active (but unmoderated) mailing list for technical discussions.

Register now to reserve a space in the Intel Developer Forum sessions on network edge technology by visiting the IDF Conference site.

Author Bio

Michael Condry recently joined Intel to spearhead its Network Edge effort as director of Network Edge Technologies. He came to the company from Sun Microsystems, where he created development processes, ran the Internet standards organization, and was heavily involved in quality management. Michael also led a research group in High Performance Network technology at the University of Illinois, where he taught computer science. Previously, he worked in VLSI architecture and UNIX* research at AT&T Bell Labs. Michael chairs a working group for the Internet Engineering Task Force. He is also on the board of directors for both the Broadband Content Delivery Forum and for the IEEE Industrial Electronics Society. Michael received his B.A. in mathematics from West Virginia University, and his M.S. and Ph.D. in computer science from Yale University.

—End of Intel Developer Update Magazine Issue 23—